Opening Persistent systems to the Internet

RT3.2 Task Force

September 10, 1998

Through the World-Wide-Web (abbr. Web) service, everybody having an Internet connection has now access to huge amount of data, structured through hyperlinks. However this mine of data still lacks many features to make it a convenient repository. It lacks the features currently available in PAS and notably high level languages for interrogation or updates.

Obviously, there is a need for the PAS and Web worlds to collaborate. Programmers developing large PAS applications will want to utilize, in addition to the data stored in their system, information available on the Web. Conversely, information stored in some PAS should be made available to the Web users.

Our objectives are to make this collaboration possible. As a first step, we have designed and implemented a tool to allow the communication from/to a PAS to/from the Web. The prototype is based on ideas discussed during the PASTEL worshops with the various partners of Task RT3.2 and has been developed at INRIA.

1 Introducing the YAT System

The number of intra/extra/internet applications integrating heterogeneous data is rapidly increasing. Central to these applications is the conversion of data from one format to another. Yet, data conversion is usually done in an *ad hoc* manner, by developing non reusable softwares. We believe that this problem requires a declarative approach and present the YAT system that achieves this goal.

To illustrate the problem, consider the following application. A car dealer company wants to build an intranet application. Among other things, the company stores informations about its dealers in a relational system and descriptions of the cars it sells in SGML documents. The idea is to integrate all information in an object PAS and to provide an HTML interface so that employees can view it on the Web. Using the YAT system that has been developed at INRIA, see Figure 1, the application programmer first imports two generic programs providing an ODMG view of relational and SGML data. Using the graphical interface, the programs are combined and customized for this specific application. Then, an ODMG to HTML program is imported and customized in a similar way. If the application requires a different Web display of one particular car or class of cars, the program can be further customized to provide different behaviours. Finally, all programs can be type checked so as to verify their coherence and composed if the required integration is virtual (i.e., the relational and SGML data are not migrated to the object system).

YAT relies on a data model allowing a uniform representation of any data, a rule based language and an instantiation mechanism allowing customization, type checking programs combination and composition.

The Model: In a manner similar to, e.g., [PGMW95, BDS95, FFLS97, ACM97], we use a simple representation of graphs to capture data coming from heterogeneous sources. The YAT model is based on named ordered trees. Its main originality is its ability to capture various levels of representation. YAT can be instantiated into other models, either ground (i.e., real data) or not (e.g., grammar/model, meta schema or schema). This property offers many advantages: it simplifies the specification and customization of conversion/integration programs, allows type checking to be turned on or off, programs to be composed or combined in a coherent way.

The Language: YATL is rule-based. It features enhanced pattern matching facilities, powerful restructuration primitives and, as in [Mai86, HY90, KKS92, PAGM96, FFLS97], explicit skolem functions. Its main originality is its ability to preserve or reconstruct the order of a collection and to perform grouping on collections featuring duplicates. Both these features are needed by most standard formats (e.g., ODMG, HTML, SGML).

The Instantiation Mechanism: Instantiation of programs and models is the fundamental characteristic of the YAT system. It provides the ability to re-use, type, combine or compose programs.

No matter the user-friendliness of a language or graphical interface, most programmers do not like to start from scratch. The common behaviour consists in fetching bits and pieces of programs and to adapt them. With YAT, this is done automatically and in a coherent way. The programmer starts with a general program that can be instantiated into one that corresponds to the input data before being customized. For instance, through successive instantiations/customizations, we may find programs generating HTML pages from (i) any input, (ii) ODMG compliant data, (iii) data corresponding to a specific ODMG schema or (iv) an object of that schema.

YAT provides flexible typing. Programs signature are infered and can be checked against any input or output formats. In this way, the programmer can verify, for instance, that the output of a program is indeed an instance of the ODMG model or some ODMG schema. Typing can be performed before or at run time.

By combination, one can easily write programs with heterogeneous inputs or that mix general and specific behaviours. This again relies on the instantiation mechanism that allows to check the coherency of a combination of programs and, if needed, to order the program rules in an appropriate manner.

If the output of a program is an instance of the input of another, they can be composed. Program composition provides the means to avoid unecessary conversion into some intermediate formats. E.g., if we consider the example application, we can combine programs so as to convert relational and SGML data directly into HTML.

The prototype: The YAT prototype implements all the features presented here. It is programmed in Objective CAML [], JAVA [] is used for the graphical interface. An on-line demonstration is available[1].

There exist several data conversion/integration languages (e.g., [DH84, AH88, Kos96, ACM97, PGMU96, PAGM96]) as well as many query languages for semi-structured data offering similar capabilities (e.g., []). However, none offers the typing flexibility of YAT or the ability to deal in an appropriate manner with collections such as lists or bags.

[1] Surf the YAT system homepage at `http://www-rocq.inria.fr/verso/Jerome.Simeon/YAT/`, for a description of the system, up-to-date implementation status, and on-line demonstration.

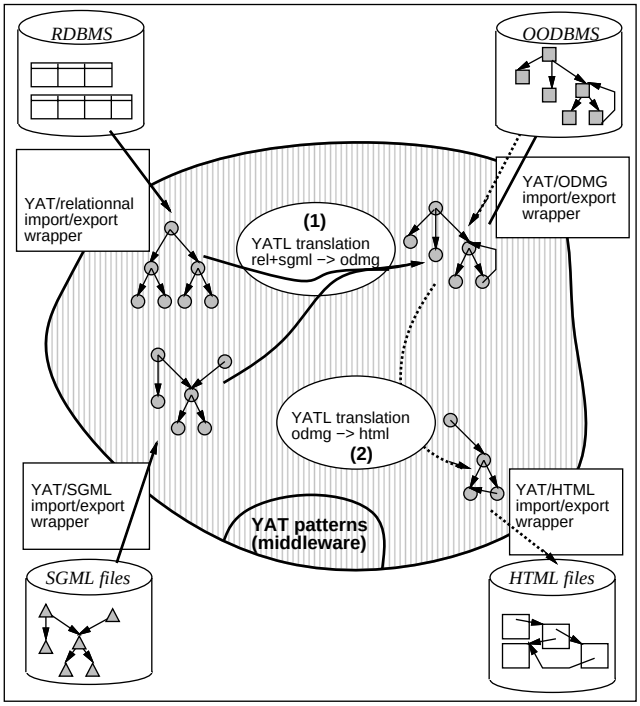

Figure 1: Translation Scenario

The sequel of the paper is organized as follows. In Section 2, we introduce the model and instantiation mechanism. The language is presented through some examples in Section 3. Section 4 shows how one can use instantiation to customize, combine and compose programs. Section 5 presents the YAT prototype and Section 6 concludes the paper.

2 The YAT data model

Since we are interested in modeling data coming from arbitrary heterogeneous sources, we need a flexible representation model. The current trend in the literature on data integration is to propose graph or tree models (e.g., [Abi97, BDS95, PGMW95, ACM97]). The main reason for the popularity of this kind of model is its simplicity and the fact that one can easily map anything into a tree/graph. Thus, and unsurprisingly, we propose yet another tree data model. We represent data in a way that does not differ much from that of [ACM97].

Yet, there is a difference. Whereas, previous work proposed models whose instances were data trees/graphs, we propose a model that can be instantiated (or refined) into another. This property is the main originality of the YAT model, and we will see that it is indeed an essential feature.

Let us illustrate this using the example scenario introduced in the previous section. We concentrate on the representation of ODMG data. A formal definition of the model and the instantiation mechanism can be found in [CS97]. Figure 2 shows four models — two of which incompletely drawn — that illustrate different levels of data representation. At the top left side of the figure, the Yat model captures any data. On its right, a first instance model of Yat captures ODMG compliant data. The Car Schema model on the bottom left side is an instance of both previous models and represents data that complies to a specific ODMG schema. Finally, the Golf model represents the actual database (we only represented the single object $\boldsymbol{c1}$) and is

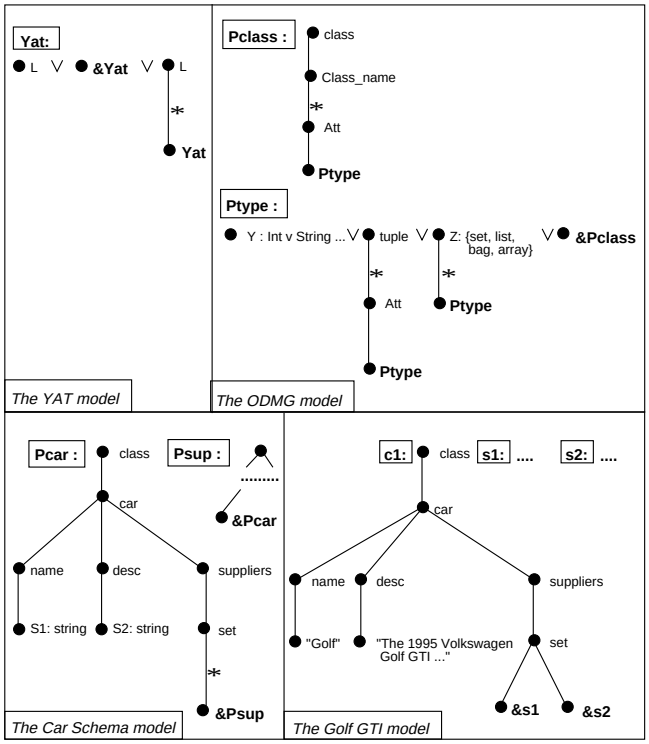

Figure 2: From General to Specific

an instance of all its predecessors.

A Model consists of a set of patterns and their associated variables domains. We start by defining variables domains.

- We use uppercase letters to denote variables. There exist two kinds of variables: (i) data variables (e.g., L in YAT, $Class\_name$ in ODMG, $S1$ in Car Schema) and (ii) pattern variables (e.g., $\boldsymbol{Yat}$, $\boldsymbol{Pclass}$, $\boldsymbol{Pcar}$). We use bold fonts for patterns.

  By default, the domain of a data variable is the set of all data constants and variable names. For instance, the domain of the Yat variable L includes constants such as $class$, car or "Golf" and variables such as $Class\_name$ and Att. Note that constants can be either symbols (e.g., $class$, $name$) or atomic data (e.g., "Golf", 1995). The default domain can be restricted. For instance, consider the ODMG model and its $\boldsymbol{Ptype}$ pattern. The domain of Y is defined as the union of $string$, int, etc. Thus, a variable whose domain is a string or an actual string like "Golf" are part of its domain.

  The domain of a pattern variable is the set of all its instance patterns. In this example, the domain of $\boldsymbol{Yat}$ includes $\boldsymbol{c1}$ but also $\boldsymbol{Pcar}$ and $\boldsymbol{Pclass}$.

- A pattern represents a structure. It is identified by a name (that can be constant or variable) and is defined by a *union* of pattern trees. A pattern tree is an ordered tree whose nodes are labeled with data variables or constants. Additionally, leaves may also be labeled with pattern names (e.g., $\boldsymbol{Yat}$, $\boldsymbol{Ptype}$) or references to pattern names (e.g., $\boldsymbol{\&Yat}$, $\boldsymbol{\&Pclass}$, $\boldsymbol{\&s1}$). The former denotes dereferencing (i.e., a pattern name label will be instantiated by a pattern tree), and is used to represent deeply recursive tree structures. The latter resembles references to objects in a standard object model and allows sharing and the management of cyclic structures (as in the Car Schema model).

The edges of a pattern tree are labeled with indicators of occurrences that are used in the instantiation process. In the example, there are two such indicators: "$\star$" which indicates zero or more occurrences and the empty symbol which indicates exactly one occurrence.

A pattern whose value is not a union (i.e. is defined by a single pattern tree), that does not contain any variable and whose edges are all empty labeled is called a *ground* pattern. E.g., ***c1*** is a ground pattern but ***Yat*** is not. A ground pattern can only be instantiated by itself. Ground patterns are used to represent real data, like in usual semistructured data models.

Model Instantiation relies on pattern instantiation which itself relies on variable domain inclusion. More precisely, (i) each pattern of the instance model must be an instance of some pattern of the source model and (ii) a variable can be instantiated either by a constant belonging to the variable's domain or by a variable whose domain is a subset.

Let us illustrate pattern instantiation by first considering why ***Pclass*** is an instance of ***Yat***. We have explained variable instantiation. Edges instantiation relies on indicators of occurrence. An empty labeled edge can only be replaced by a similar edge. A "$\star$" labeled edge can be replaced by any ordered sequence of edges, with or without label. Thus, assuming that ***Ptype*** is an instance of ***Yat*** (which is indeed the case) and considering the rightmost argument of the Yat union, we can see that the subtree whose root is *Att* is an instance of ***Yat***. The same can be said of the subtree whose root is *Class_name* (we use pattern dereferencing here). Going up one more step, we see that ***Pclass*** is indeed an instance of ***Yat***.

Now, consider the ***Pclass*** and ***c1*** patterns. The variable *Class_name* has been instantiated by the *car* symbol. The following "$\star$" edge is instantiated by a sequence of empty labeled edges. Consider the leftmost one. It goes to a node whose label *name* is an instance of the variable *Att*. The next edge is correctly empty labeled and goes to the node labeled "Golf". "Golf" is an instance of ***Ptype*** since it is a string and thus belongs to the domain of the Y variable.

To conclude this presentation of the model, let us point out the fact that there is a clear analogy between the instantiation process and subtyping. Indeed, through instantiation (resp. subtyping), we obtain patterns (resp. types) whose domains are more and more restricted. Therefore, the YAT model combines some features of typed data models, with the flexibility required to manage heterogeneous data. This property of the YAT model can be used for various purposes: to help the programmer write conversion programs or type them (see Section 3), to customize, combine or compose programs (see Section 4).

Finally, we introduce the YAT syntax for patterns with an example. Patterns ***Pcar*** and ***Psup*** from Figure 2 are given below. A labeled node is denoted by its label followed by the list of its sons between brackets($\langle$ ... $\rangle$), edges are denoted by arrows with their corresponding label. If a node has a single son, brackets can be omitted.

Pattern for Car Objects

$$\begin{aligned} \boldsymbol{Pcar} \ : \ & class \rightarrow car\langle \rightarrow name \rightarrow S1 : String, \\ & \quad \rightarrow desc \rightarrow S2 : String, \\ & \quad \rightarrow suppliers \rightarrow set \xrightarrow{\star} \boldsymbol{\&Psup} \ \rangle \end{aligned}$$

Pattern for Supplier Objects

$$\boldsymbol{Psup} : class \rightarrow supplier\langle \rightarrow name \rightarrow S1 : String, \\ \rightarrow city \rightarrow S2 : String, \\ \rightarrow zip \rightarrow S3 : String\rangle$$

3 The YAT language

In this section, we present YATL, our language for the specification of data conversion. This language supports a graphical interface. Thus, programmers do not have to write YATL programs, they are generated by the system given a graphical specification.

We introduce YATL basic features, demonstrate the language ability to handle heterogeneity problems between data sources and present advanced YATL features to deal with collections. Then, we address the problem of detecting cycles in a program. Finally, we show how YATL programs can be type checked. A formal definition can be found in [CS97].

3.1 Basic features

YATL is a rule-based language, each rule describing a part of the data conversion. A rule is composed of a body and a head. The body contains patterns and boolean predicates whose roles are to *filter* the input data, and external functions that *compute* additional data. The head of a rule contains a single pattern which describes how the data which has been filtered in the body must be re-structured.

Let us illustrate this with a first program generating car and supplier objects from a set of SGML brochures. The brochures comply to the following Document Type Definition (i.e. its grammar or DTD) and give information about the cars, such as year of construction, technical description and suppliers. The number element is distinct for each brochure, and the title element stores the name of the car.

Brochures DTD

```
<!DOCTYPE   brochure [
<!ELEMENT   brochure   - -  (number, title,
                             model, desc, spplrs) >
<!ELEMENT   number     - -  (#PCADATA) >
<!ELEMENT   title      - -  (#PCADATA) >
<!ELEMENT   model      - -  (#PCADATA) >
<!ELEMENT   desc       - -  (#PCADATA) >
<!ELEMENT   spplrs     - -  (sup)* >
<!ELEMENT   supplier   - -  (name,address) >
<!ELEMENT   name       - -  (#PCADATA) >
<!ELEMENT   address    - -  (#PCADATA) > ] >
```

The first rule of the program creates the supplier objects. Figure 3 illustrates sample input and output for Rule 1.

Rule 1

$$
\begin{array}{ll}
 & \boldsymbol{Pbr}\ : \\
 & brochure\langle \rightarrow number \rightarrow Num, \\
\boldsymbol{Psup}(SN)\ : & \quad \rightarrow title \rightarrow T, \\
class \rightarrow supplier & \quad \rightarrow model \rightarrow Year, \\
\quad \langle \rightarrow name \rightarrow SN, \Longleftarrow & \quad \rightarrow desc \rightarrow D, \\
\quad \rightarrow city \rightarrow C, & \quad \rightarrow suppls \stackrel{\star}{\rightarrow} supplier \\
\quad \rightarrow zip \rightarrow Z\rangle & \quad\quad \langle \rightarrow name \rightarrow SN, \\
 & \quad\quad\quad \rightarrow address \rightarrow Add\rangle\rangle, \\
 & Year > 1975, \\
 & C \texttt{ is } \mathrm{city}(Add), \\
 & Z \texttt{ is } \mathrm{zip}(Add)
\end{array}
$$

Note that we have not given a domain to the various data variables used in the rule. This means that they are associated to the default domain and can be instantiated by any constant. The body of the rule is on the right-hand side of the $\Longleftarrow$ symbol. It consists of (i) a pattern representing the SGML brochures, (ii) one simple predicate which eliminates oldest cars and (iii) two external functions which extract city and zip code from an input address. The head of the rule consists of a single pattern whose name is parameterized by a variable (SN) representing the name of a supplier. Parameterized pattern names correspond to explicit skolem functions as found in [KKS92, PAGM96, FFLS97]. Intuitively, the above skolem states that one object will be created for each supplier name encountered in the input SGML brochures. Thus, if the same supplier name appears twice in the input, only one object will be created.

This somehow implies that the name of a supplier represents its key. If it is not the case, the result of transformation described by Rule 1 is non deterministic (e.g. if two suppliers share the same name SN and not the same address, then YATL semantics do not specify the value chosen for $\boldsymbol{Psup}(SN)$). Syntactical restrictions can guarantee the determinism of a program. However, this has a cost that the user may not want to pay (e.g., the ability to remove duplicate values). In YAT, we accept potentially non-deterministic programs and alert the user at run time when the same pattern name is associated to two distinct values.

Let us now use the above example to explain how rules are applied on input patterns **b1**, **b2** from Figure 3. The input of a rule is always a set of ground patterns (i.e., completely instantiated patterns) which is processed in five phases.

1. Each pattern in the input set is matched against the body of the rule thus forming the following set of variable bindings.
   { [$\boldsymbol{Pbr} = \boldsymbol{b1}$; $Num = 1$; $T =$ "Golf"; ...; $SN =$ "VW center"; $Add =$ "Bd Lenoir ..."]
   [$\boldsymbol{Pbr} = \boldsymbol{b2}$; $Num = 2$; $T =$ "Golf"; ...; $SN =$ "VW2"; $Add =$ "Bd Leblanc ..."]
   [$\boldsymbol{Pbr} = \boldsymbol{b2}$; $Num = 2$; $T =$ Golf"; ... ; $SN =$ "VW Center"; $Add =$ "Bd Lenoir ..."] }

2. External functions are evaluated to generate new bindings. The first element of the above set then becomes:
   [$\boldsymbol{Pbr} = \boldsymbol{b1}$; $Num = 1$; $T =$ "Golf"; ...; $Add =$ "Bd Lenoir ..."; $C =$ "Paris"; $Z =$ 75005]

   External functions are typed. This means that a type filter is applied on the set of variable bindings before they are evaluated.

3. Predicates are applied to filter the set of variable bindings. In our case, all bindings are kept.

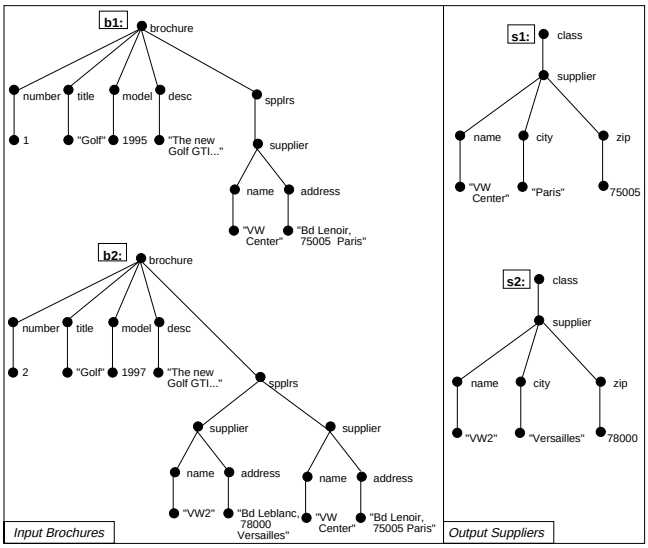

Figure 3: Applying Rule 1 on two SGML brochures

4. Skolem functions are evaluated. In this case, $\boldsymbol{s1}$ is associated to $\boldsymbol{Psup}$("VW center") and $\boldsymbol{s2}$ is associated to $\boldsymbol{Psup}$("VW2"). Skolem functions are not dependent of a given rule but are global to a program. In other words, the association between $\boldsymbol{Psup}$("VW center") and $\boldsymbol{s1}$ can be recalled in another rule of the program.

5. Finally, the set of output patterns is constructed (in a straightforward manner as far as this example is concerned) and each output pattern is associated to its name (see Figure 3).

Note that step 4 and 5 can be mixed since it is not necessary to evaluate all skolem functions before computing the corresponding pattern tree. Let us now consider the second and last rule of our program. It creates one car object for each brochure, each car referencing its set of suppliers.

Rule 2

$$
\begin{array}{l}
\boldsymbol{Pcar}(\boldsymbol{Pbr}) : \\
class \rightarrow car \\
\quad \langle \rightarrow name \rightarrow T, \\
\quad\ \rightarrow desc \rightarrow D, \\
\quad\ \rightarrow suppliers \\
\qquad \rightarrow set \overset{\{\}}{\rightarrow} \boldsymbol{\&Psup}(SN) \rangle
\end{array}
\quad \Longleftarrow \quad
\begin{array}{l}
\boldsymbol{Pbr} : \\
brochure \\
\quad \langle \rightarrow number \rightarrow Num, \\
\quad\ \rightarrow title \rightarrow T, \\
\quad\ \rightarrow model \rightarrow Year, \\
\quad\ \rightarrow desc \rightarrow D, \\
\quad\ \rightarrow suppls \overset{\star}{\rightarrow} supplier \\
\qquad \langle \rightarrow name \rightarrow SN, \\
\qquad\ \rightarrow address \rightarrow Add \rangle \rangle
\end{array}
$$

The creation of links from cars to suppliers is handled by the use of the parameterized pattern name $\boldsymbol{Psup}(SN)$. Because skolem functions are global to a program and can be processed independently from their values, Rule 1 and Rule 2 can be applied in any order. So, it is not necessary to use stratification mechanism to preserve YATL declarativity. Note that we use here the **&** symbol to denote a reference to a supplier. In other words, the corresponding leaves in the output pattern tree will point to a supplier pattern tree (e.g., $\boldsymbol{\&s1}$). Omitting the **&**

symbol entails a dereferenciation, e.g., ***s1*** would be replaced by its associated tree. Remark that it requires that the value associated to ***s1*** exists. Therefore, dereferenciation is handled at the end of rules processing.

Note the use of symbol {} in the above rule's head. This annotation is used to create a node (here, the node *set*) with several sons, all distinct and in no specified order. In the above rule, node *set* will aggregate all the references to the suppliers of a given car.

Let us now see how cyclic references are managed in YATL. For this, we redefine the way suppliers are created. Rule 1′ shows how cars can be associated to their suppliers.

Rule 1′

$$
\begin{array}{l}
\boldsymbol{Psup}(SN) : \\
class \rightarrow supplier \\
\quad \langle \rightarrow name \rightarrow SN, \\
\quad\ \rightarrow city \rightarrow C, \\
\quad\ \rightarrow zip \rightarrow Z, \\
\quad\ \rightarrow sells \\
\quad\quad \rightarrow set \xrightarrow{\{\}} \boldsymbol{\&Pcar}(\boldsymbol{Pbr})\rangle
\end{array}
\quad \Longleftarrow \quad
\begin{array}{l}
\boldsymbol{Pbr} : \\
brochure \\
\quad \langle \rightarrow number \rightarrow Num, \\
\quad\ \rightarrow title \rightarrow T, \\
\quad\ \rightarrow model \rightarrow Year, \\
\quad\ \rightarrow desc \rightarrow D, \\
\quad\ \rightarrow suppls \xrightarrow{\star} supplier \\
\quad\quad \langle \rightarrow name \rightarrow SN, \\
\quad\quad\quad \rightarrow address \rightarrow Add\rangle\rangle, \\
C \texttt{ is city}(Add), \\
Z \texttt{ is zip}(Add)
\end{array}
$$

Remember that the **&** symbol before a pattern name corresponds to a reference. Should we remove the **&** symbol in both Rules 1′ and 2, we would face a cyclic program. Detection of cyclic programs is addressed in Section 3.4.

3.2 Dealing with heterogeneity

To illustrate YATL's ability to handle heterogeneity between multiple data sources, we propose to write a rule taking its data from two distinct sources: a relational database whose schema is given below and the above SGML documents.

Suppliers Relational Database Schema

suppliers[sid: integer, name: string, city: string,
address: string, tel: string]
cars[cid: integer, broch_num: string]
sales[sid : integer, cid: integer, year: integer, sold: integer]

Rule 3 creates one car object for each distinct car in the relational database if it corresponds to a brochure. Note that the SN variable is used in both body patterns to indicate that the supplier name in the relational database and in the SGML documents should be the same. The external function 'sameaddress' takes care of the heterogeneity between addresses in the database and in the SGML documents.

Rule 3

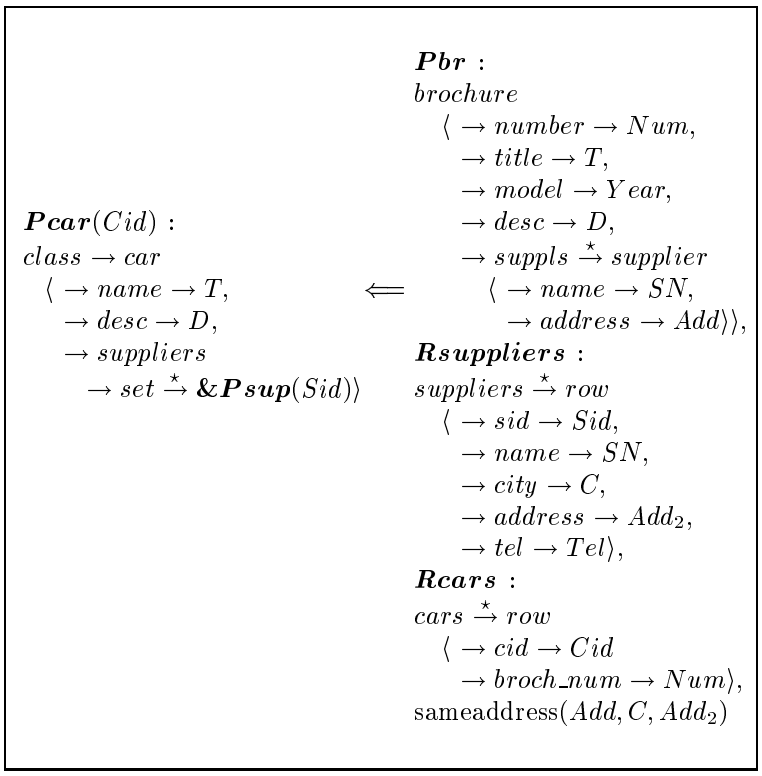

This rule shows how one can handle eventual inconsistencies in the data sources using the YAT environment. Note that a similar approach to handle heterogeneity has been followed in [ACM97, PGMU96, PAGM96].

3.3 Dealing with collections

Like other languages for data conversion/integration, YATL handles sets. However, an original feature of the language is that, due to its data model and its active domain semantics (as opposed to fix-point), it is also well adapted to the manipulation of collections such as multi-sets and lists. YATL provides one primitive for grouping and another for ordering. These two simple features are sufficient to deal appropriately with collections. We illustrate this point with two examples.

Rule 4 creates an ODMG list of supplier objects (as created by Rule 1') ordered by their name. For this it uses a primitive that combines both grouping (to remove duplicates) and ordering on a given criterion. This requires that the set of variable bindings be processed (in that case grouped and ordered) before the actual construction of the corresponding edges. In the example, the criterion is a single variable. More variables can be used. Also, one can perform grouping on some criteria and ordering on some others.

Rule 4

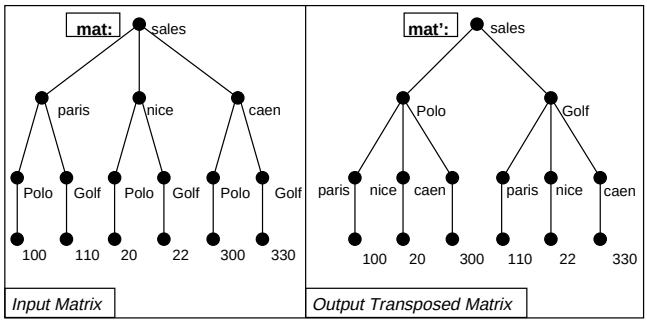

Figure 4: Transposing a 3×2 Matrix

$$\begin{array}{lll} & & \boldsymbol{Pbr}\ : \\ & & brochure \\ & & \quad \langle \rightarrow number \rightarrow Num, \\ \boldsymbol{Pallsups}()\ : & & \quad \rightarrow title \rightarrow T, \\ list \stackrel{[]_{SN}}{\rightarrow} \boldsymbol{\&Psup}(SN) & \Longleftarrow & \quad \rightarrow model \rightarrow Year, \\ & & \quad \rightarrow desc \rightarrow D, \\ & & \quad \rightarrow suppls \stackrel{\star}{\rightarrow} supplier \\ & & \qquad \langle \rightarrow name \rightarrow SN, \\ & & \qquad \rightarrow address \rightarrow Add \rangle \rangle \end{array}$$

The possibility to handle ordered collections is important for various application domains. Ordered collections exist in many standard data models, and are essential to deal with documents. The above example was rather database oriented. The YATL language is also able to manipulate arrays, which are the subject of increasing interest [FM95, LMW96, MS97]. This is illustrated by the following rule that transposes any input matrix. Note the use of the index edges I and J to store the information about the original ordering of the sons.

Rule 5

$$\begin{array}{lll} \boldsymbol{New(Id)}\ : & & \boldsymbol{Id}\ : \\ Mat \stackrel{[]_J}{\rightarrow} Y \stackrel{[]_I}{\rightarrow} X \rightarrow A & \Longleftarrow & Mat \stackrel{I}{\rightarrow} X \stackrel{J}{\rightarrow} Y \rightarrow A \end{array}$$

Figure 4 illustrates the application of Rule 5 on a matrix containing statistics on car sales.

3.4 Detecting cyclic programs

Remember our comment about the program presented in Section 3.1 which consisted of Rules 1' and 2: if we remove the **&** symbol in both rules we potentially introduce a cycle in the program. Indeed, let us consider a car $\boldsymbol{c1}$ and a supplier $\boldsymbol{s1}$ referencing each other. By removing the **&** symbol, we break the referencing mechanism and require that the $\boldsymbol{s1}$ tree includes the $\boldsymbol{c1}$ tree which itself should include the $\boldsymbol{c1}$ tree, etc.

The bad news is that the general problem of statically detecting cycles in a YATL program is undecidable (it is similar to finding wrong recursive calls in a programming language). However,

we may detect possibly cyclic programs and forbid their use. One simple way to do that is to construct the dependency graph of dereferenced Skolems. If this graph is cyclic, it implies a potential cycle in the program that is thus rejected. If we consider the above example, there exists a cycle between the ***Psup*** and ***Pcar*** skolem functors.

However, this method has a serious drawback. It rejects conversion programs working by recursion on the input tree, programs commonly used for arbitrarily deep recursive structures. Fortunately, this important class of programs features an interesting characteristic: they usually rely on a skolem functor whose sole parameter is a pattern name and whose recursive invocation is always performed on a subtree of the input. This is further explained in the next section. These recursive programs are called *safe-recursive*. This can be checked syntactically and guarantees the absence of cycles at run time.

To summarize, we reject programs that are not *safe-recursive* and that feature a cycle in the graph of dependencies of the dereferenciated Skolems.

3.5 Typing in YATL

To conclude this section, we briefly explain how typing can be performed in YATL.

A program takes a model as input and returns another model (see Section 2). Input and output models can easily be inferred by considering the program (i) input and output patterns, (ii) predicate/function signatures and (iii) variable domains.

For instance, consider the program consisting of Rule 1 and Rule 2. The input model of the program consists of the single brochure pattern ***Pbr*** with restrictions on the type of variables *Add* and *Year*. The type of *Add* is given by the signature of functions *city* and *zip*, that of *Year* by the "`>`" predicate. The output model consists of two patterns ***Pcar*** and ***Psup*** and restrictions on variables *C*, *Z*.

The couple of input/output models will be called the signature of a conversion program and noted $M_{IN} \mapsto M_{OUT}$. Now, once the signature of a program has been inferred, this information can be used in several ways.

It can be used to check if two programs can safely be composed or combined. This is explained in the next section. It can also be used to check that the input (or output) model of a program is indeed an instance of a more general model. For instance, the user may check that a program generates *car* and *supplier* objects compliant with a given ODMG schema or, more generally, with the ODMG model. Finally, by turning typing on at run time, we may verify that all input data is involved in the conversion process. This is simply done by adding a rule such as Rule Exception to the program. The rule should be applied only when another rule cannot (see the following section to understand how this is possible). It features an empty head. Its body contains one pattern that matches any input data and an external function that invokes an exception.

Rule Exception

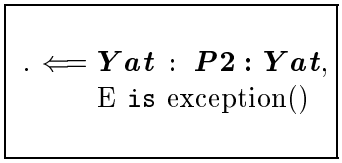

It is important to understand that typing in YAT is in no way constraining. Programs do not need it to be executed. If a part of the input data does not match the input model of a program, no conversion will be performed on it, but no error will occur.

4 Customizing, combining and composing conversion programs

One of the most interesting feature of a YATL program is its ability to be "instantiated". Program instantiation constitutes the basis of program customization. We show now how it works using an example. We next show how programs can be combined so as to provide both customized and general behaviors. Finally, we explain program composition.

4.1 Customizing programs

No matter the user-friendliness of a language or graphical interface, most programmers do not like to start from scratch. The common behavior consists in fetching bits and pieces of programs and to adapt them. With YAT, the programmer starts with a general program that can be instantiated into one that corresponds to its input data. Customization can follow. For instance, through successive instantiations/customizations, we may find programs generating HTML pages from (i) any input, (ii) ODMG compliant data, (iii) data corresponding to a specific ODMG schema or (iv) an object of that schema. Let us illustrate this with an example.

Below is a program that converts any ODMG data into HTML, thus allowing users of our Intranet application to view data from a Web browser. The specified translation is that implemented by the O_2Web system [O2W96]. An object is converted into an HTML page (Rule Web1), an atomic value into a string (Rule Web2), a collection or a tuple into a list of HTML items (Rules Web3 to Web5) and an object reference to an HTML anchor (Rule Web6). The rules bodies contain ODMG patterns, while rules heads contains HTML patterns as illustrated in Figure 5. Four facts are noteworthy.

1. The program creates a new identifier for each HTML page through the ***HtmlPage*** skolem function. It is the HTML wrapper's responsibility to map these pattern identifiers to a real URL when creating the actual HTML pages.
2. The program uses recursive dereferenciation on skolem ***HtmlElement*** (in rules Web3, Web4 and Web5), but is *safe-recursive* since the skolem parameter is a pattern name (***Ptype***) and the recursive calls are performed on a subtree ($\boldsymbol{P_2}$).
3. Some edges are labeled with the $\star$ symbol. This label provides implicit grouping as the $\{\}$ symbol but without duplicate elimination.
4. Finally, many variables are typed. For instance, X in Rule Web4 (resp. Web5) can only be instantiated with the *set* or *bag* (resp. *list* or *array*) symbol. $\boldsymbol{P_2}$ in Rules Web3 to Web6 can only be instantiated by a pattern name instance of ***Ptype***.

Web1

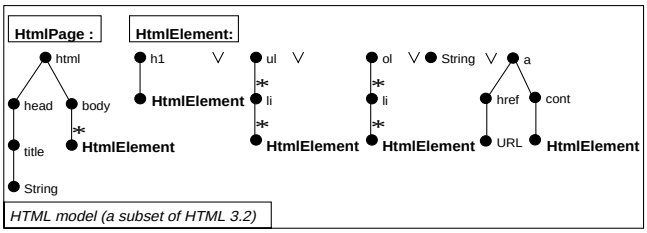

Figure 5: HTML patterns

$$
\begin{array}{l}
\boldsymbol{HtmlPage}(\boldsymbol{Pclass}) : \\
html\langle \rightarrow head \rightarrow title \rightarrow Classname, \\
\quad \rightarrow body\langle \rightarrow h1 \rightarrow Classname, \\
\quad\quad \rightarrow ul \xrightarrow{\star} li \\
\quad\quad\quad \langle \rightarrow Att2, \\
\quad\quad\quad\quad \rightarrow \boldsymbol{HtmlElement}(\boldsymbol{Ptype})\rangle\rangle\rangle \\
\quad \Longleftarrow \\
\quad\quad \boldsymbol{Pclass} : \\
\quad\quad class \rightarrow Classname \xrightarrow{\star} Att \rightarrow \boldsymbol{Ptype}, \\
\quad\quad Att2 \texttt{ is } \mathrm{concat}(Att, " : ")
\end{array}
$$

Web2

$$
\begin{array}{l}
\boldsymbol{HtmlElement}(\boldsymbol{Ptype}) : S \Longleftarrow \boldsymbol{Ptype} : Data, \\
\qquad\qquad\qquad\qquad\qquad\qquad S \texttt{ is } \mathrm{data\_to\_string}(Data)
\end{array}
$$

Web3

$$
\begin{array}{l}
\boldsymbol{HtmlElement}(\boldsymbol{Ptype}) : \\
ul \xrightarrow{\star} li\langle \rightarrow Att2, \\
\quad \rightarrow \boldsymbol{HtmlElement}(\boldsymbol{P_2 : Ptype})\rangle \\
\quad \Longleftarrow \\
\quad\quad \boldsymbol{Ptype} : \\
\quad\quad tuple \xrightarrow{\star} Att \rightarrow \boldsymbol{P_2 : Ptype} \\
\quad\quad Att2 \texttt{ is } \mathrm{concat}(Att, " : ")
\end{array}
$$

Web4

$$
\begin{array}{l}
\boldsymbol{HtmlElement}(\boldsymbol{Ptype}) : \\
ul \xrightarrow{\star} li \rightarrow \boldsymbol{HtmlElement}(\boldsymbol{P_2 : Ptype})\rangle \\
\quad \Longleftarrow \\
\quad\quad \boldsymbol{Ptype} : \\
\quad\quad X : \{set, bag\} \xrightarrow{\star} \boldsymbol{P_2 : Ptype}
\end{array}
$$

Web5

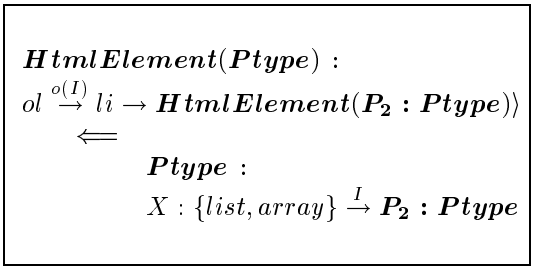

Web6

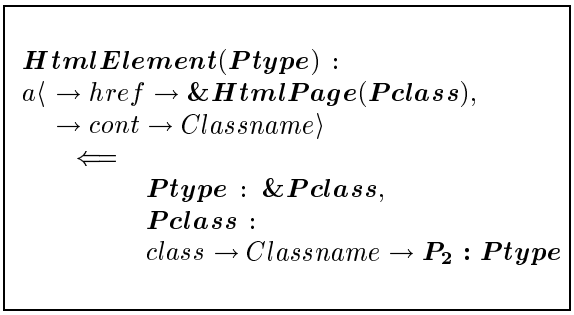

Now, let us suppose that we want to display objects of the class car in a different way and the Golf GTI object in yet another way. Using a tool such as O_2Web, this would entail writing appropriate methods from scratch using C++ or some other O_2 supported programming language. With YAT, the user instantiates the general program by giving a more specific pattern. This instantiation process is done automatically, and the resulting new program is equivalent to the previous one, but more specific. This program can then be customized accordingly to the programmer's needs. Let us illustrate this by instantiating the Web program through the ***Pcar*** pattern introduced in Section 3 (see Rule 2). Instantiation results in the following Rule.

rule WebCar

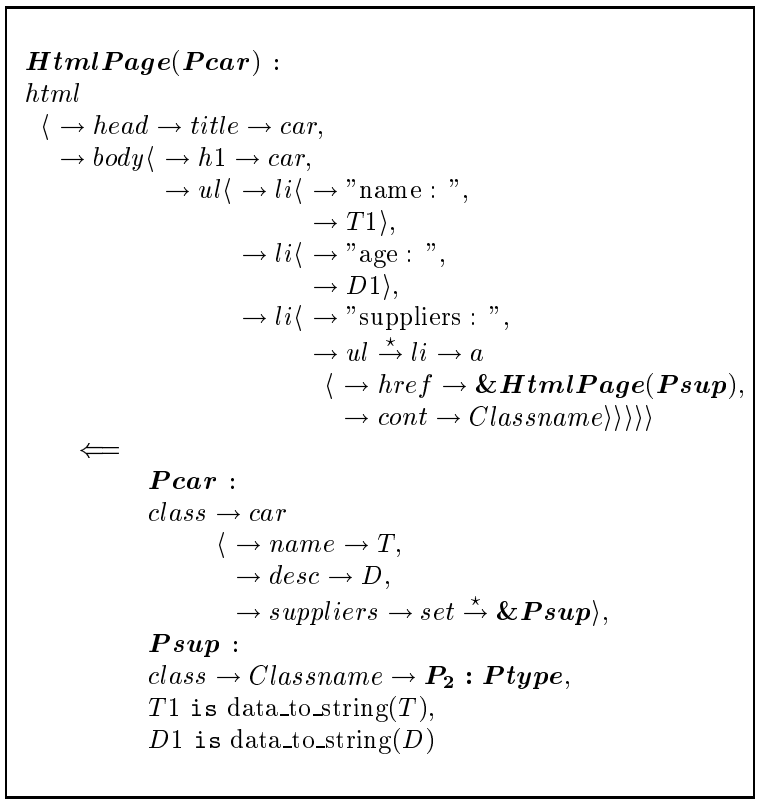
HtmlPage(Pcar) :
html
⟨ → head → title → car,
→ body⟨ → h1 → car,
→ ul⟨ → li⟨ → "name : ",
→ T1⟩,
→ li⟨ → "age : ",
→ D1⟩,
→ li⟨ → "suppliers : ",
→ ul ⋆→ li → a
⟨ → href → &HtmlPage(Psup),
→ cont → Classname⟩⟩⟩⟩
⇐
Pcar :
class → car
⟨ → name → T,
→ desc → D,
→ suppliers → set ⋆→ &Psup⟩,
Psup :
class → Classname → P2 : Ptype,
T1 is data_to_string(T),
D1 is data_to_string(D)

The body of the rule contains two function calls and two patterns: one is the given ***Pcar*** pattern, the other one is an incomplete ***Psup*** pattern[2] which has been obtained through instantiation of rule Web6. It specifies that pattern name ***Psup*** can only be instantiated by an object belonging to some class.

Let us see how Rule WebCar has been derived. Four rules have been used: Rule Web1 on the upper part of the car pattern, Rule Web2 on the two subtrees associated to properties *name* and *desc*, Rule Web4 on the subtree associated to the *supplier* property and Rule Web6 on the leaf containing a reference to ***Psup***. Then, output trees have been generated and, due to the dereferenciation, have been appended together to form the head part of the rule. One exception is the pattern that matched the ***&Psup*** leaf. Since it has not been dereferenced, it has been added to the rule body along with all encountered function calls. Remark that two calls to the 'data_to_string' external functions are derived from rule Web2 and therefore, the system must provide appropriate renaming of variables (here to $T1$ and $D1$) to avoid conflicts.

Once derived, the new rule can be rewritten to provide a customized presentation for car objects. For instance, if we do not want to display the supplier of a car, we can rewrite Rule WebCar in the following way:

rule newWebCar

[2]Since we only give the ***Pcar*** pattern for the instantiation, the system does not assume any knowledge of the ***Psup*** pattern referenced in ***Pcar***.

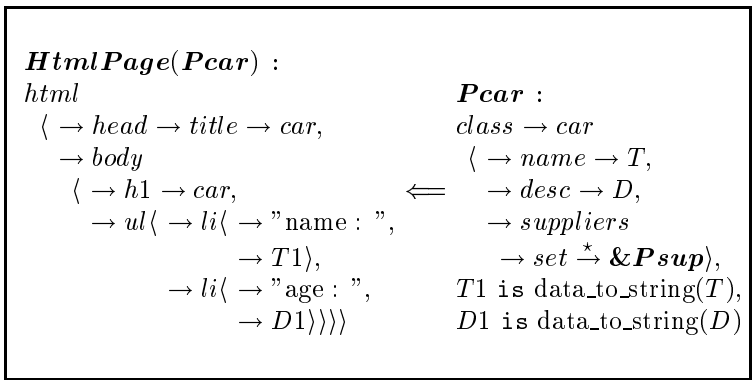

4.2 Combining programs

Rule WebCar can be considered as the unique rule of a program translating car objects. Obviously, we would like to combine this program with the previous one so as to be able to convert all objects into appropriate HTML pages. However, by doing this without caution we would provide two translations for a given car object.

To solve this problem, YATL interpreter organizes the set of rules of a program hierarchically. Thus, a program now consists of a set of rule hierarchies. For a given input pattern, the more specific rules (leaves in the hierarchy) matching the input are applied first. If matching cannot be obtained, less specific rules in the hierarchy are tried and so on.

Rule hierarchies are built by the YATL interpreter according to possible rule conflicts. A conflict occurs only when: (i) there is a subtype relationship between two rules input models (as for rules WebCar and Web1 in the above example), and (ii) the skolem functions used in these rules are the same (e.g. there is no conflict for rules 1 and 2 defined in Section 3 as they do not code for the same set of output patterns).

Finally, the user may also enforce the rule hierarchy to apply some rules in a particular order, according to the conversion he wants to achieve. For that purpose, YATL allows to enforce a specific hierarchy of rules. Of course, in this case, the declarativity of YATL programs is transgressed. As was noted in the previous section, this ability to organize rules can be used to check programs at run time.

4.3 Composing programs

Let us now explain program composition. Recall the example scenario and the program consisting of Rules 1 and 2 that created cars and suppliers from a given brochure. In order to avoid unnecessary loading into the object database, the user may want to combine this program with the above one that performs ODMG to HTML conversion. This would provide a direct conversion from SGML documents to HTML.

A first solution would be to apply successively the two programs. However, this would result in unnecessary processing, since the system would create intermediate ODMG patterns. Using YATL, it is possible to *compose* these two programs and generate a more efficient program.

Let us explain the principle of the composition mechanism. Taking two conversion programs $prg1 : M_1 \mapsto M_2$ whose input (resp. output) model is M_1 (resp. M_2) and $prg2 : M_2' \mapsto M_3$, the system first check if $prg1$ and $prg2$ are compatible (i.e. if M_2 is an instance of M_2'). If this is the case, the system instantiates $prg2$ with the patterns of M_2. This produces a new

translation $prg2' : M_2 \mapsto M_3'$. Then, the final composition is straightforward as syntactically identical patterns appear in the output model of $prg1$ and the input model of $prg2'$.

The signature of the two programs that we want to combine (SGML to ODMG and ODMG to HTML) are compatible. Therefore, the ODMG to HTML translation can be instantiated on ***Pcar*** and ***Psup*** patterns. It generates a rule WebCar$'$ (which only differs from Rule WebCar by having additional informations about pattern ***Psup***) and another rule WebSup. Finally, the program containing Rule 1 and Rule 2 is composed with this instantiated program. Two new rules result from this composition. The first, from the composition of Rule 1 and Rule WebSup, creates HTML pages for the suppliers contained in our SGML brochures. The second, from the composition of Rule 2 and Rule WebCar$'$, generates HTML pages for the car descriptions. It is given below.

Rule (2+Webcar$'$)

$$
\begin{array}{l}
\boldsymbol{HtmlPage}(\boldsymbol{Pcar}) : \\
html \\
\quad \langle \rightarrow head \rightarrow title \rightarrow car, \\
\quad\ \rightarrow body\langle \rightarrow h1 \rightarrow car, \\
\qquad\qquad \rightarrow ul\langle \rightarrow li\langle \rightarrow \text{"name : "}, \\
\qquad\qquad\qquad\qquad \rightarrow T1\rangle, \\
\qquad\qquad\qquad \rightarrow li\langle \rightarrow \text{"age : "}, \\
\qquad\qquad\qquad\qquad \rightarrow D1\rangle, \\
\qquad\qquad\qquad \rightarrow li\langle \rightarrow \text{"suppliers : "}, \\
\qquad\qquad\qquad\qquad \rightarrow ul \xrightarrow{\star} li \rightarrow a \\
\qquad\qquad\qquad\qquad\quad \langle \rightarrow href \rightarrow \&\boldsymbol{HtmlPage}(SN), \\
\qquad\qquad\qquad\qquad\qquad \rightarrow cont \rightarrow supplier\rangle\rangle\rangle\rangle \\
\quad \Longleftarrow \\
\qquad\quad \boldsymbol{Pbr} : \\
\qquad\quad brochure\langle \rightarrow number \rightarrow Num, \\
\qquad\qquad\qquad\quad \rightarrow title \rightarrow T, \\
\qquad\qquad\qquad\quad \rightarrow model \rightarrow Year, \\
\qquad\qquad\qquad\quad \rightarrow desc \rightarrow D, \\
\qquad\qquad\qquad\quad \rightarrow suppls \xrightarrow{\star} supplier \\
\qquad\qquad\qquad\qquad \langle \rightarrow name \rightarrow SN, \\
\qquad\qquad\qquad\qquad\quad \rightarrow address \rightarrow Add\rangle\rangle, \\
\qquad\quad T1 \texttt{ is } \text{data\_to\_string}(T), \\
\qquad\quad D1 \texttt{ is } \text{data\_to\_string}(D)
\end{array}
$$

5 Architecture and implementation

In this section, we briefly present the architecture of the YAT System as well as the current status of the prototype.

5.1 Architecture

The YAT system provides a complete environment to design data conversion programs. It does not provide the ability to query or update external sources. Still, we believe that it can serve as

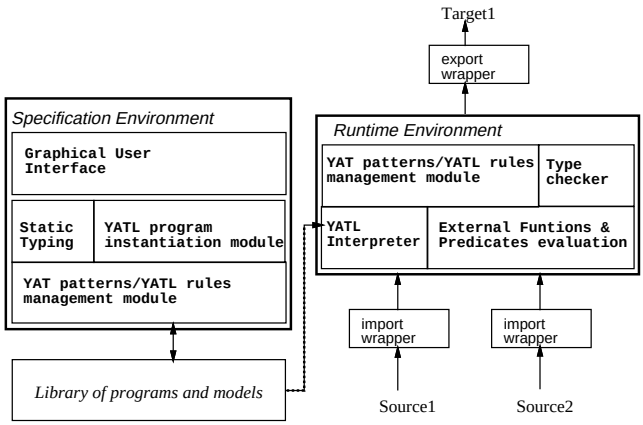

Figure 6: YAT System architecture

the basis for a mediator/wrapper system for data integration and are currently working in that direction.

The heart of the system is composed of: (i) a module to manipulate YAT data and programs, (ii) a module for type checking and type inference that can be called *on demand* and (iii) a YATL interpreter.

Figure 6 describes the current architecture. There are three main parts : (i) the specification environment, (ii) the run-time environment and (iii) a library of programs and formats. Both environments rely on the "YAT model/YATL rules management module" for the internal representation and manipulation of patterns and rules. The specification environment provides (i) a module to statically check/infer the type of a program and (ii) a graphical user interface to specify translations, using (iii) program instantiation to customize existing programs if needed. The execution environment uses wrappers to import (resp. export) data into (resp. from) YAT and an interpreter to perform the translation. The interpreter relies on a separate module to process external functions and predicates and, if required by the user, a type checker. The library allows to save and import programs and models.

The runtime environment can be used independently or be linked to import/export wrappers to generate stand-alone executables (e.g. like LaTeX2HTML). In this case the translation system connects to the import program to retrieve the input data whenever required. If the HTML output wrapper is used, the generated executable can be used as a CGI script.

5.2 Prototype

The YAT system is implemented in the Verso Database Group at INRIA. It is written in Objective CAML [Ler97] and JAVA (for the graphical interface). It is now fully operational. It provides all the above features as well as some export/import wrappers (HTML, O_2 database and XML) and appropriate conversion programs. Figures 7 and 8 show the YATL editor main window with Rule $1'$ and the pattern editor. [Sma97] gives a complete description of this graphical interface.

We are currently enhancing this first version of the prototype with new wrappers, and plan to work on performance issues.

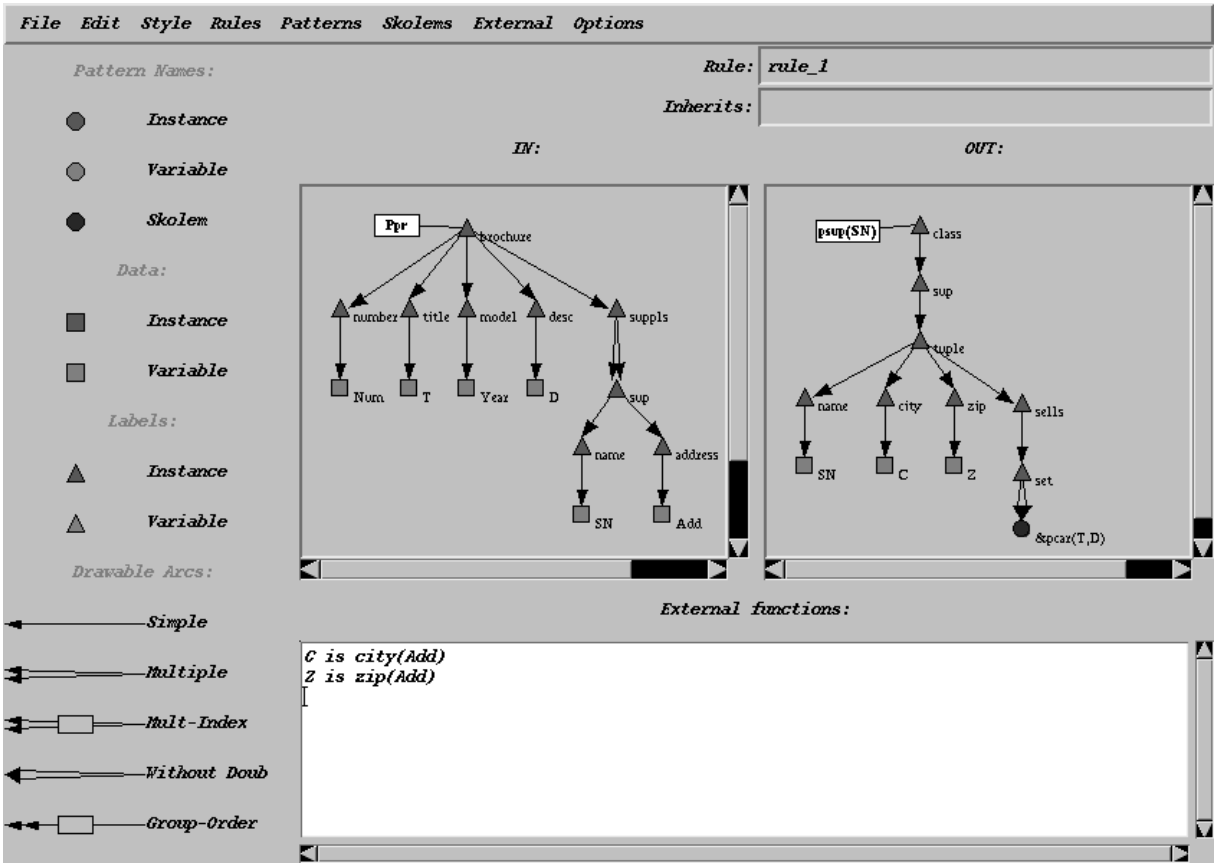

Figure 7: Rule 1′ using YATL Editor

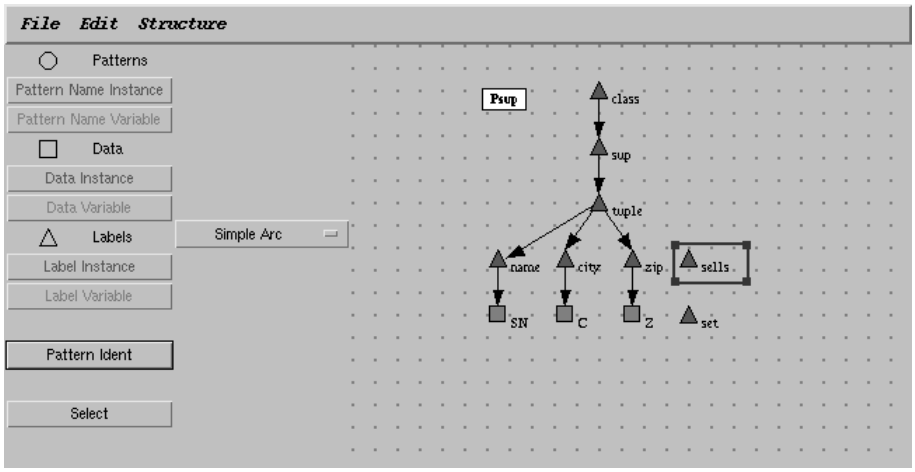

Figure 8: YATL Pattern Editor

6 Conclusion

So far, we have focused on data conversion. The next step would be to provide a system for data integration, i.e., provide the means to queryexternal data as well as convert it.However, some issues should be addressed to achieve this goal. Notably, we plan to work on the following issues:

Design and implementation of a query language : YATL lacks essential features to be an adequate query language. Notably, one should be able to formulate keyword queries. Furthermore, we need to be able to translate queries from YAT to the sources that are queried. This obviously requires smarter wrappers and mediator.

Management of persistency : Although the mediator can be viewed as a simple gateway, it should be able to store some data so as to be able to answer fast to a query or to palliate to source server crashes.

Incremental view maintainance : As soon as data is stored in the system, we must be able to refresh it in an efficient way and when needed.

These three directions of research will be further investigate with the various partners of task RT3.2.

References

[Abi97] S. Abiteboul. Querying Semi-Structured Data. In *Proceedings of International Conference on Database Theory (ICDT)*, pages 1–18, Delphi, Greece, January 1997.

[ACM97] S. Abiteboul, S. Cluet, and T. Milo. Correspondence and translation for heterogeneous data. In *Proceedings of International Conference on Database Theory (ICDT)*, Delphi, Greece, January 1997.

[AH88] S. Abiteboul and R. Hull. Restructuring hierarchical database objects. *Theoretical Computer Science*, 62:3–38, 1988.

[BDS95] P. Buneman, S. B. Davidson, and D. Suciu. Programming constructs for unstructured data. In *Proceedings of International Workshop on Database Programming Languages*, Gubbio, Italy, 1995.

[CS97] S. Cluet and J. Siméon. Data integration based on data conversion and restructuring. Technical report, Verso database group - INRIA, October 1997. `http://www-rocq.inria.fr/verso/Jerome.Simeon/YAT/`.

[DH84] U. Dayal and H. Hwang. View definition and generalisation for database integration in multibase: A system for heterogeneous distributed databases. *IEEE Transactions on Software Engineering*, 10(6):628–644, November 1984.

[FFLS97] M. F. Fernandez, D. Florescu, A. Y. Levy, and D. Suciu. A query language for a web-site management system. *SIGMOD Record*, 26(3):4–11, September 1997.

[FM95] L. Fegaras and D. Maier. Towards an effective calculus for object query languages. In *Proceedings of ACM SIGMOD Conference on Management of Data*, pages 47–58, San Jose, California, May 1995.

[HY90] R. Hull and M. Yoshikawa. Ilog: Declarative creation and manipulation of object identifiers. In *Proceedings of International Conference on Very Large Databases (VLDB)*, pages 455–468, Brisbane, Australia, August 1990.

[KKS92] M. Kifer, W. Kim, and Y. Sagiv. Querying object-oriented databases. In *Proceedings of ACM SIGMOD Conference on Management of Data*, pages 393–402, San Diego, California, June 1992.

[Kos96] A. S. Kosky. *Transforming Databases with Recursive Data Structures.* PhD thesis, University of Pennsylvania, 1996.

[Ler97] X. Leroy. *The Objective Caml system release 1.07.* INRIA, December 1997. Documentation and user's manual. `ftp://ftp.inria.fr/lang/caml-light/`.

[LMW96] L. Libkin, R. Machlin, and L. Wong. A query language for multidimensional arrays: Design, implementation, and optimization techniques. In *Proceedings of ACM SIGMOD Conference on Management of Data*, pages 229–239, Montreal, Canada, June 1996.

[Mai86] D. Maier. A logic for objects. In *Workshop on Foundations of Deductive Databases and Logic Programming*, pages 6–26, Washington, D.C., August 1986.

[MS97] A. P. Marathe and K. Salem. A language for manipulating arrays. In *Proceedings of International Conference on Very Large Databases (VLDB)*, pages 46–55, Athens, Greece, August 1997.

[O2W96] O_2 Technology, Versailles. *The O_2Web Reference Manual version 4.6*, September 1996.

[PAGM96] Y. Papakonstantinou, S. Abiteboul, and H. Garcia-Molina. Object fusion in mediator systems. In *Proceedings of International Conference on Very Large Databases (VLDB)*, pages 413–424, Bombay, India, September 1996.

[PGMU96] Y. Papakonstantinou, H. Garcia-Molina, and J. D. Ullman. Medmaker: A mediation system based on declarative specifications. In *Proceedings of IEEE International Conference on Data Engineering (ICDE)*, pages 132–141, New Orleans, Louisiana, February 1996.

[PGMW95] Y. Papakonstantinou, H. Garcia-Molina, and J. Widom. Object exchange across heterogeneous information sources. In *Proceedings of IEEE International Conference on Data Engineering (ICDE)*, pages 251–260, Taipei, Taiwan, March 1995.

[Sma97] K. Smaga. Interface graphique pour la traduction de données. Master's thesis, Université Paris VI-CNAM-ENST, September 1997.